arXiv:1608.08899v1 [cs.CG] 31 Aug 2016

Visibility Representations of Boxes in 2.5 Dimensions*

Alessio Arleo[1], Carla Binucci[1], Emilio Di Giacomo[1], William S. Evans[2], Luca Grilli[1], Giuseppe Liotta[1], Henk Meijer[3], Fabrizio Montecchiani[1], Sue Whitesides[4], and Stephen Wismath[5]

[1] Università degli Studi di Perugia, Italy, alessio.arleo@studenti.unipg.it
{carla.binucci, emilio.digiacomo, luca.grilli, giuseppe.liotta, fabrizio.montecchiani}@unipg.it
[2] University of British Columbia, Canada, will@cs.ubc.ca
[3] University College Roosevelt, the Netherlands, h.meijer@ucr.nl
[4] University of Victoria, Canada, sue@uvic.ca
[5] University of Lethbridge, Canada, wismath@uleth.ca

Abstract. We initiate the study of 2.5D box visibility representations (2.5D-BR) where vertices are mapped to 3D boxes having the bottom face in the plane $z = 0$ and edges are unobstructed lines of sight parallel to the x- or y-axis. We prove that: (*i*) Every complete bipartite graph admits a 2.5D-BR; (*ii*) The complete graph K_n admits a 2.5D-BR if and only if $n \leqslant 19$; (*iii*) Every graph with pathwidth at most 7 admits a 2.5D-BR, which can be computed in linear time. We then turn our attention to 2.5D grid box representations (2.5D-GBR) which are 2.5D-BRs such that the bottom face of every box is a unit square at integer coordinates. We show that an n-vertex graph that admits a 2.5D-GBR has at most $4n - 6\sqrt{n}$ edges and this bound is tight. Finally, we prove that deciding whether a given graph G admits a 2.5D-GBR with a given footprint is NP-complete. The footprint of a 2.5D-BR Γ is the set of bottom faces of the boxes in Γ.

1 Introduction

A *visibility representation (VR)* of a graph G maps the vertices of G to non-overlapping geometric objects and the edges of G to *visibilities*, i.e., segments that do not intersect any geometric object other than at their end-points. Depending on the type of geometric objects representing the vertices and on the rules used for the visibilities, different types of representations have been studied in computational geometry and graph drawing.

A *bar visibility representation (BVR)* maps the vertices to horizontal segments, called *bars*, while visibilities are vertical segments. BVRs were introduced in the 80s as a modeling tool for VLSI problems [18, 29, 30, 36–38]. The graphs

* Research started at the 2016 Bertinoro workshop on Graph Drawing. Research supported in part by NSERC, and by MIUR project AMANDA prot. 2012C4E3KT_001.

that admit a BVR are planar and they have been characterized under various models [18, 30, 36, 38].

Extensions and generalizations of BVRs have been proposed in order to enlarge the family of representable graphs. In a *rectangle visibility representation (RVR)* the vertices are axis-aligned rectangles, while visibilities are both horizontal or vertical segments [4, 7, 12, 14, 15, 25, 31, 33]. RVRs can exist only for graphs with thickness at most two and with at most $6n - 20$ edges [25]. Recognizing these graphs is NP-hard in general [31] and can be done in polynomial time in some restricted cases [4, 33]. Generalizations of RVRs where orthogonal shapes other than rectangles are used to represent the vertices have been recently proposed [17, 28]. Another generalization of BVRs are *bar k-visibility representations (k-BVRs)*, where each visibility segment can "see" through at most k bars. Dean et al. [13] proved that the graphs admitting a 1-BVR have at most $6n - 20$ edges. Felsner and Massow [22] showed that there exist graphs with a 1-BVR whose thickness is three. The relationship between 1-BVRs and 1-planar graphs has also been investigated [1, 9, 19, 34].

RVRs are extended to 3D space by *Z-parallel Visibility Representations (ZPR)*, where vertices are axis-aligned rectangles belonging to planes parallel to the xy-plane, while visibilities are parallel to the z-axis. Bose et al. [8] proved that K_{22} admits a ZPR, while K_{56} does not. Štola [32] subsequently reduced the upper bound on the size of the largest representable complete graph by showing that K_{51} does not admits a ZPR. Fekete et al. [20] showed that K_7 is the largest complete graph that admits a ZPR if unit squares are used to represent the vertices. A different extension of RVRs to 3D space are the *box visibility representations (BR)* where vertices are 3D boxes, while visibilities are parallel to the x-, y- and z- axis. This model was studied by Fekete and Meijer [21] who proved that K_{56} admits a BR, while K_{184} does not.

In this paper we introduce *2.5D box visibility representations (2.5D-BR)* where vertices are 3D boxes whose bottom faces lie in the plane $z = 0$ and visibilities are parallel to the x- and y-axis. Like the other 3D models that use the third dimension, 2.5D-BRs overcome some limitations of the 2D models. For example, graphs with arbitrary thickness can be realized. In addition 2.5D-BRs seem to be simpler than other 3D models from a visual complexity point of view and have the advantage that they can be physically realized, for example by 3D printers or by using physical boxes. Furthermore, this type of representation can be used to model visibility between buildings of a urban area [11]. The main results of this paper are as follows.

- We show that every complete bipartite graph admits a 2.5D-BR (Section 3). This implies that there exist graphs that admit a 2.5D-BR and have arbitrary thickness.
- We prove that the complete graph K_n admits a 2.5D-BR if and only if $n \leqslant 19$ (Section 3). Thus, every graph with $n \leqslant 19$ vertices admits a 2.5D-BR.
- We describe a technique to construct a 2.5D-BR of every graph with pathwidth at most 7, which can be computed in linear time (Section 4).

– We then study *2.5D grid box representations* (*2.5D-GBR*) which are 2.5D-BRs such that the bottom face of every box is a unit square with corners at integer coordinates (Section 5). We show that an n-vertex graph that admits a 2.5D-GBR has at most $4n - 6\sqrt{n}$ edges and that this bound is tight. It is worth remarking that VRs where vertices are represented with a limited number of shapes have been investigated in the various models of visibility representations. Examples of these shape-restricted VRs are unit bar VRs [16], unit square VRs [12], and unit box VRs [21].
– Finally, we prove that deciding whether a given graph G admits a 2.5D-GBR with a given footprint is NP-complete (Section 5). The *footprint* of a 2.5D-BR Γ is the set of bottom faces of the boxes in Γ.

For reasons of space, some proofs and details are omitted and can be found in the appendix.

2 Preliminaries

A *box* is a six-sided polyhedron of non-zero volume with axis-aligned sides in a 3D Cartesian coordinate system. In a *2.5D box representation* (*2.5D-BR*) the vertices are mapped to boxes that lie in the non-negative half space $z \geqslant 0$ and include one face in the plane $z = 0$, while each edge is mapped to a *visibility* (i.e. a segment whose endpoints lie in faces of distinct boxes and whose interior does not intersect any box) parallel to the x- or to the y-axis. We remark that visibilities between non-adjacent objects may exist, i.e., we adopt the so called *weak visibility model* (in the *strong visibility model* each visibility between two geometric objects corresponds to an edge of the graph). The weak model seems to be the most effective when representing non-planar graphs and it has been adopted in several works (see e.g. [4, 9, 19]). As in many papers on visibility representations [21, 26, 33, 35, 38], we assume the *ϵ-visibility model*, where each segment representing an edge is the axis of a positive-volume cylinder that intersects no box except at its ends; this implies that an intersection point between a visibility and a box belongs to the interior of a box face. In what follows, when this leads to no confusion, we shall use the term *edge* to indicate both an edge and the corresponding visibility, and the term *vertex* for both a vertex and the corresponding geometric object.

Given a box b of a 2.5D-BR, the face that lies in the plane $z = 0$ is called the *footprint* of b. The intersection of the plane $z = 0$ with a 2.5D-BR Γ is called the *footprint* of Γ and is denoted by Γ_0. In other words, the footprint of a 2.5D-BR Γ consists of the footprint of all the boxes in Γ. If Γ is a 2.5D-BR of a complete graph then its footprint Γ_0 satisfies a trivial necessary condition (throughout the paper we will refer to this condition as *NC*): for every pair of boxes b_1 and b_2 of Γ, there must exist a line ℓ (in the plane $z = 0$) such that (i) ℓ passes through the footprints of b_1 and b_2, and (ii) ℓ is either parallel to the x-axis or to the y-axis. A *2.5D grid box representation* (*2.5D-GBR*) is a 2.5D-BR such that every box has a footprint that is a unit square with corners at integer coordinates.

Two boxes *see each other* if there exists a visibility between them; we say that they *see each other above another box* b, if there exists a visibility between them and the projection of this visibility on the plane $z = 0$ intersects the interior of the footprint of b. Notice that this implies that the two boxes are both taller than b. We say that two boxes have a *ground visibility* or are *ground visible* if there exists a visibility between their footprints, i.e. if there exists an unobstructed axis-aligned line segment connecting their footprints. If two boxes are ground visible then they see each other regardless of their heights and the heights of the other boxes. Let G be a graph, let Λ be a collection of boxes each lying in the non-negative half space $z \geqslant 0$ with one face in the plane $z = 0$, such that the boxes of Λ are in bijection with the vertices of G. Note that Λ may not be a 2.5D-BR of G. For a vertex v of G, $\Lambda(v)$ denotes the corresponding box in Λ, while $h(\Lambda(v))$, or simply $h(v)$, indicates the height of this box. For a subset $S \subset V(G)$, $\Lambda(S)$ denotes the subset of boxes associated with S, while $\Lambda_0(S)$ is the footprint of $\Lambda(S)$. Let $G[S]$ be the subgraph of G induced by S. We say that $\Lambda(S)$ is a 2.5D-BR of $G[S]$ in Λ, if for any edge (u, v) of $G[S]$ there exists a visibility in Λ between $\Lambda(u)$ and $\Lambda(v)$; that is, the visibility is not destroyed by the presence of the other boxes in Λ.

3 2.5D Box Representations of Complete Graphs

In this section we consider 2.5D-BRs of complete graphs and complete bipartite graphs.

Theorem 1. *Every complete bipartite graph admits a 2.5D-BR.*

Proof. Let $K_{m,n}$ be a complete bipartite graph. We represent the m vertices in the first partite set with m boxes $a_0, a_1, \ldots, a_{m-1}$ such that box a_i has a footprint with corners at $(2i, 0, 0)$, $(2i + 1, 0, 0)$, $(2i, 2n - 1, 0)$ and $(2i + 1, 2n - 1, 0)$ and height $m - i$. Then we represent the n vertices in the second partite set with n boxes $b_0, b_1, \ldots, b_{n-1}$ such that box b_j has a footprint with corners at $(2m, 2j, 0)$, $(2m + 1, 2j, 0)$, $(2m, 2j + 1, 0)$ and $(2m + 1, 2j + 1, 0)$ and height m. Consider now a box a_i and a box b_j. By construction a_i and b_j see each other above all boxes a_l with $l > i$. □

A consequence of Theorem 1 is that there exist graphs with unbounded thickness that admit a 2.5D-BR. This contrasts with other models of visibility representations (e.g., k-BVRs, and RVRs), which can only represent graphs with bounded thickness.

We now prove that the largest complete graph that admits a 2.5D-BR is K_{19}. We first show that given a 2.5D-BR of a complete graph there is one line parallel to the x-axis and one line parallel to the y-axis whose union intersect all boxes and such that each of them intersects at most 10 boxes. This implies that there can be at most 20 boxes in a 2.5D-BR of a complete graph. We then show that there must be a box that is intersected by both lines, thus lowering this bound to 19. We finally exhibit a 2.5D-BR of K_{19}. We start with some technical lemmas. The proof of the next one can be found in Appendix A.

Lemma 1. *Let G be an n-vertex graph that admits a 2.5D-BR Γ'. Then there exists a 2.5D-BR Γ of G such that every box of Γ has a distinct integer height in the range $[1, n]$ and the footprint of Γ is the same as that of Γ'.*

The following lemma is proved in [27, Obervation 1]; we give an alternative proof in Appendix A. Given an axis-aligned rectangle r in the plane $z = 0$, we denote by $x(r)$ the x-extent of r and by $y(r)$ the y-extent of r, so $r = x(r) \times y(r)$.

Lemma 2. [27] *For every arrangement $\mathcal{R}$ of n axis-aligned rectangles in the plane such that for all $a, b \in \mathcal{R}$, either $x(a) \cap x(b) \neq \emptyset$ or $y(a) \cap y(b) \neq \emptyset$, there exists a vertical and a horizontal line whose union intersects all rectangles in $\mathcal{R}$.*

The following lemma is similar to the Erdős–Szekeres lemma and can be proved in a similar manner [20]. A sequence of distinct integers is *unimaximal* if no element of the sequence is smaller than both its predecessor and successor.

Lemma 3. [20] *For all $m > 1$, in every sequence of $\binom{m}{2} + 1$ distinct integers, there exists at least one unimaximal sequence of length m.*

Given a 2.5D-BR Γ and a line ℓ parallel to the x-axis or to the y-axis, we say that ℓ *stabs* a set of boxes B of Γ if it intersects the interior of the footprints of each box in B. Let $b_1, b_2, \ldots, b_h$ be the boxes of B in the order they are stabbed by ℓ. We say that B has a *staircase layout*, if $h(b_i) > h(b_{i-1})$ for $i = 2, 3, \ldots, h$.

Lemma 4. *In a 2.5D-BR of a complete graph no line parallel to the x-axis or to the y-axis can stab five boxes whose heights, in the order in which the boxes are stabbed, form a unimaximal sequence.*

Proof. Assume, as a contradiction, that there exists a line ℓ parallel to the x-axis or to the y-axis that stabs 5 boxes $b_1, \ldots, b_5$ whose heights form a unimaximal sequence in the order in which the boxes are stabbed by ℓ. Let r_i be the footprint of box b_i (with $1 \leqslant i \leqslant 5$). We claim that there exists a ground visibility between every pair of boxes b_i and b_j (with $1 \leqslant i < j \leqslant 5$). If $j = i+1$ this is clearly true. Suppose then that $j \neq i + 1$. If b_i and b_j do not have a ground visibility, then they must see each other above b_l with $i < l < j$, i.e., the height of b_i and of b_j must be larger than the height of b_l, which is impossible because the sequence of heights is unimaximal. Thus, for every pair of boxes b_i and b_j there must be a ground visibility. Since b_i and b_j are both stabbed by ℓ, this visibility must be parallel to ℓ. This implies that the left sides (if ℓ is parallel to the x-axis) or the bottom sides (if ℓ is parallel to the y-axis) of rectangles r_1, r_2, r_3, r_4, r_5 form a bar visibility representation of K_5, which is impossible because bar visibility representations exist only for planar graphs [23]. □

Lemma 5. *In a 2.5D-BR of a complete graph no line parallel to the x-axis or to the y-axis can stab more than* 10 *boxes.*

Proof. Let Γ be a 2.5D-BR of a complete graph K_n. By Lemma 1 we can assume that all boxes have distinct integer heights. Suppose, as a contradiction, that

there exists a line ℓ parallel to the x-axis or to the y-axis that stabs $k > 10$ boxes. Let $h_1, h_2, \ldots, h_k$ be the heights of the stabbed boxes in the order in which the boxes are stabbed by ℓ. By Lemma 3 this sequence of heights contains a unimaximal sequence of length 5, but this is impossible by Lemma 4. □

Lemma 6. *A complete graph admits a 2.5D-BR only if it has at most* 19 *vertices.*

Proof. Let Γ be a 2.5D-BR of a complete graph K_n (for some $n > 0$). By Lemma 1 we can assume that all boxes of Γ have distinct heights. The footprint Γ_0 of Γ is an arrangement of rectangles that satisfies Lemma 2. Thus there exist a line ℓ_h parallel to the x-axis and a line ℓ_v parallel to the y-axis that together stab all boxes of Γ. By Lemma 5, both ℓ_h and ℓ_v can stab at most 10 boxes each. This means that the number of boxes (and therefore the number of vertices of K_n) is at most 20. We now prove that if ℓ_h and ℓ_v both stab ten boxes, there must be one box that is stabbed by both ℓ_h and ℓ_v, which implies that the number of boxes in Γ is at most 19.

Suppose, for a contradiction, that $p = \ell_h \cap \ell_v$ does not lie in a box. Refer to Figure 1(a) for an illustration. Denote by T the set of boxes stabbed by ℓ_v that are above p and by B be the set of boxes stabbed by ℓ_v that are below p. Analogously, denote by L the set of boxes stabbed by ℓ_h that are to the left of p and by R the set of boxes stabbed by ℓ_h that are to the right of p. Each of these sets can be empty but $|T|+|B| = 10$ and $|L|+|R| = 10$. Denote by $l_1, l_2, \ldots, l_{|L|}$ the set of boxes in L from right to left, i.e., l_1 is the box closest to p. Analogously, denote by $r_1, r_2, \ldots, r_{|R|}$ the boxes of R from left to right (r_1 is the closest to p), by $t_1, t_2, \ldots, t_{|T|}$ the boxes of T from bottom to top (t_1 is the closest to p) and by $b_1, b_2, \ldots, b_{|B|}$ the boxes of B from top to bottom (b_1 is the closest to p). Let f_T, f_B, f_L, and f_R be the footprints of t_1, b_1, l_1, and r_1, respectively. Let ℓ_X be the line containing the side of f_X that is closest to p and let ℓ'_X be the line containing the opposite side of f_X (for every $X \in \{T, B, L, R\}$).

We first claim that for each f_X there exists a line ℓ_Y (with $X, Y \in \{T, B, L, R\}$ and $Y \neq X$) that intersects the interior of f_X. Suppose, for a contradiction, that this is not true for at least one f_X, say f_L; that is, the interior of f_L is not intersected by ℓ_T and ℓ_B. If so, there must be a line ℓ parallel to the y-axis that intersects all the rectangles in $T \cup B$ and f_L; otherwise the necessary condition NC does not hold for $T \cup B \cup \{l_1\}$. But then ℓ would stab eleven boxes, which is impossible by Lemma 5. Thus, our claim holds and the four rectangles f_X are placed so that ℓ_T, ℓ_R, ℓ_B, and ℓ_L stab f_R, f_B, f_L, and f_T (or, symmetrically, f_L, f_T, f_R, and f_B, which follows a symmetric argument), respectively, as in Figure 1(a).

We consider now the sets T, B, L, and R. For each set there are two possible configurations. Consider the set B and the line ℓ'_L. If the set $B' = B \setminus \{b_1\}$ contains a box b_j whose footprint is completely to the right of ℓ'_L, we say that B has *configuration A* (see Figure 1(b)). In the case of configuration A, the footprint of all boxes in $L' = L \setminus \{l_1\}$ must extend below the line ℓ'_B (otherwise the necessary condition NC does not hold for $L' \cup \{b_j\}$). This implies that $y(f_B)$

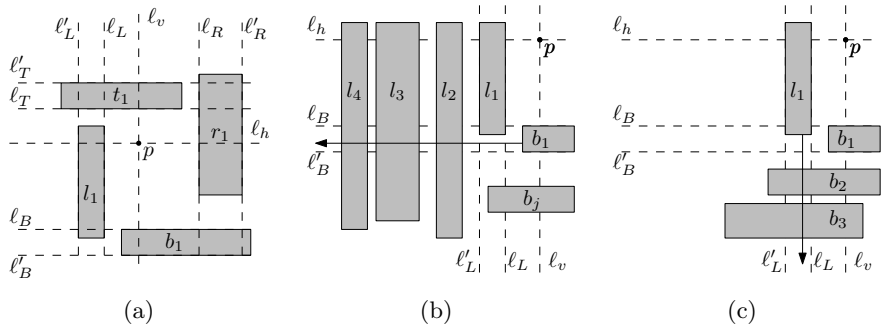

Fig. 1: (a) Placement of the four rectangles f_T, f_R, f_B, and f_L. (b) Configuration A for the boxes of set B. (c) Configuration B for the boxes of set B. The arrow intersects the boxes that must have a staircase layout.

is contained in $y(l_i)$ for all $i \geqslant 2$. The only possibility for b_1 to see all these boxes is that L' has a staircase layout (with l_2 being the shortest box) and b_1 is taller than the second tallest one. So, configuration A for the set B implies that L' has a staircase layout. If all boxes of B' have a footprint that extends to the left of ℓ'_L, we say that B has *configuration B* (see Figure 1(c)). In this case, $x(f_L)$ is contained in $x(b_i)$ for all $i \geqslant 2$. Again, the only possibility for l_1 to see all these boxes is that B' has a staircase layout and that l_1 is taller than the second tallest one. So, configuration B for the set B implies that B' has a staircase layout. The definitions of configurations A and B for T, L, R are similar to those for B and arise by considering lines ℓ'_R, ℓ'_T, ℓ'_B, respectively.

For any two sets X and Y that are consecutive in the cyclic order T, R, B, L, either X' or Y' has a staircase layout (depending on whether X has configuration A or B). This implies that either B' and T' have both a staircase layout or L' and R' have both a staircase layout. Suppose that B' and T' have a staircase layout (the case when L' and R' have a staircase layout is analogous). If either $|B'| \geqslant 5$ or $|T'| \geqslant 5$, ℓ_v stabs at least five boxes whose heights form a unimaximal sequence, which is impossible by Lemma 4. Thus $|B'| = 4$ and $|T'| = 4$ (recall that $|B'| + |T'| = 8$). Since all boxes of Γ have distinct heights, either $h(b_2) < h(t_2)$ or $h(t_2) < h(b_2)$. In the first case ℓ_v stabs the five boxes t_5, t_4, t_3, t_2, b_2 whose heights form a unimaximal sequence, which is impossible by Lemma 4. In the other case ℓ_v stabs the five boxes b_5, b_4, b_3, b_2, t_2 whose heights form a unimaximal sequence, which is impossible by Lemma 4. □

We conclude this section by exhibiting a 2.5D-BR of K_{19}, illustrated in Figure 2. To prove the correctness of the drawing the idea is to partition the vertex set of K_{19} into five subsets (shown in Figure 2) and prove that all boxes in a given set see all other boxes (details are in Appendix A). The following theorem holds.

Theorem 2. *A complete graph K_n admits a 2.5D-BR if and only if $n \leqslant 19$.*

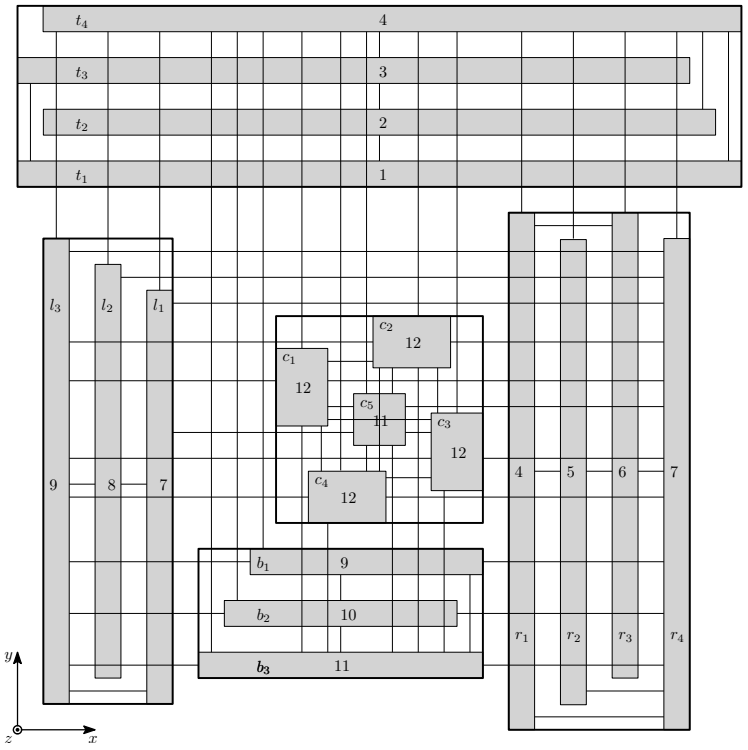

Fig. 2: Illustration of a 2.5D-BR of K_{19}, the footprint is represented by a 2D drawing in the plane $z = 0$, while the heights of boxes are indicated by integer labels. The five rectangles with thick sides represent the partitioning of $V(K_{19})$ into five subsets.

4 2.5D Box Representations of Graphs with Pathwidth at Most 7

A graph G with pathwidth p is a subgraph of a graph that can be constructed as follows. Start with the complete graph K_{p+1} and classify all its vertices as *active*. At each step, a vertex is *deactivated* and a new active vertex is introduced and joined to all the remaining active vertices. The order in which vertices are introduced is given by a *normalized path decomposition*, which can be computed in linear time for a fixed p [24]. For a definition of normalized path decomposition see Appendix B.

Theorem 3. *Every n-vertex graph with pathwidth at most 7 admits a 2.5D-BR, which can be computed in $O(n)$ time.*

Proof. We describe an algorithm to compute a 2.5D-BR of a graph G with pathwidth 7. The algorithm is based on the use of eight groups of rectangles, a

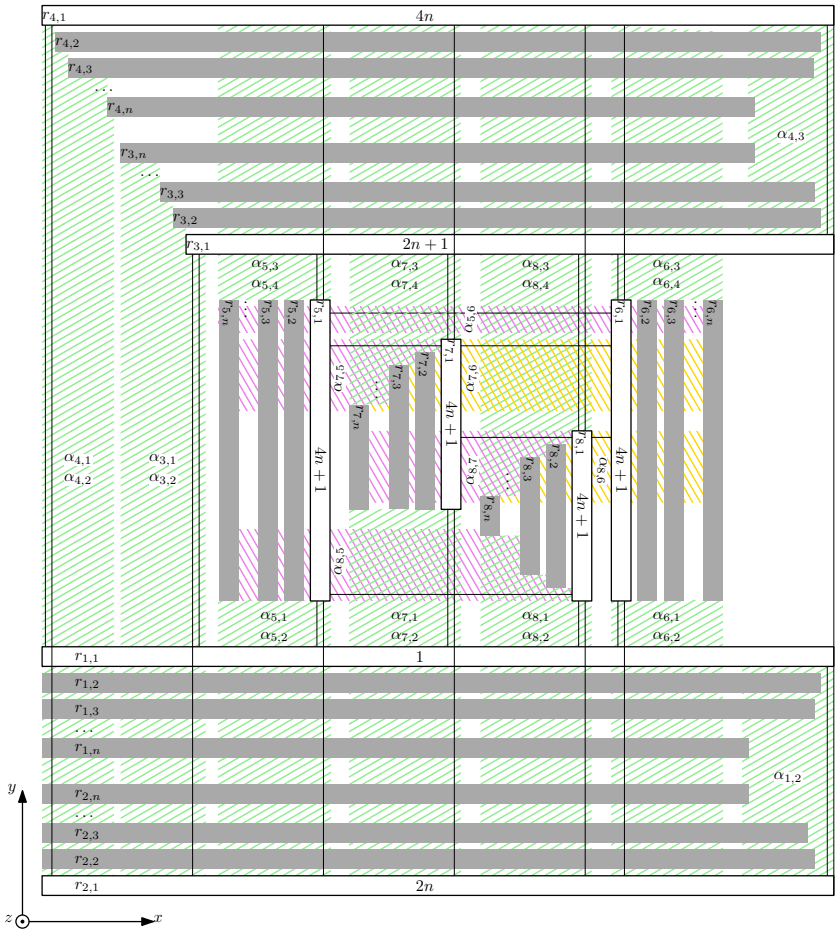

Fig. 3: Construction of a 2.5D-BR for a graph with pathwidth 7.

subset of which will form the footprint of the 2.5D-BR of G. For graphs with pathwidth $p < 7$, the same algorithm can be applied by considering only $p + 1$ groups, arbitrarily chosen.

The eight groups are defined in the plane $z = 0$ and have n rectangles each denoted as $r_{h,1}, r_{h,2}, \ldots, r_{h,n}$ $(1 \leqslant h \leqslant 8)$. The groups are placed as shown in Figure 3. The groups $h = 5, 6, 7, 8$ will be called *central groups*. A vertex whose footprint is $r_{h,k}$ will be called a *vertex of group* h $(1 \leqslant h \leqslant 8)$.

Let $v_1, v_2, \ldots, v_n$ be the vertices of G in the order given by a normalized path decomposition. We denote by G_i the subgraph of G induced by $\{v_1, v_2, \ldots, v_i\}$. We create a collection of boxes by adding one box per step; at step i we add a box to represent the next vertex v_i to be activated. We denote the collection of the first i boxes as Λ_i and we prove that Λ_i satisfies the following invariant (I1): Λ_i is a 2.5D-BR of G_i such that for any pair of boxes of group j and

k ($1 \leqslant j, k \leqslant 8$) that represent vertices that are adjacent in G_i, there exists a visibility whose projection in the plane $z = 0$ is inside the region $\alpha_{j,k}$. The regions $\alpha_{j,k}$ are highlighted in Figure 3 as dashed regions.

The initial eight active vertices $v_1, v_2, \ldots, v_8$ are represented by boxes whose footprints are $r_{1,1}, r_{2,1}, \ldots, r_{8,1}$, respectively. The heights are set as follows: $h(v_h) = (h-1) \cdot n + 1$, for $h = 1, 2, 3, 4$, and $h(v_h) = 4n + 1$ for $h = 5, 6, 7, 8$. The initial eight vertices are shown in Figure 3 as white rectangles whose heights are shown inside them. Λ_8 satisfies invariant I1 thanks to the visibilities shown in Figure 3.

Assume now that Λ_{i-1} ($i > 8$) satisfies invariant I1 and let v_j be the vertex to be deactivated (for some $j < i$). Assume that v_j belongs to group h ($1 \leqslant h \leqslant 8$). Vertex v_i is represented as a box with footprint $r_{h,i}$ and height $h(v_i) = h(v_j) + 1$, if $h \in \{1, 3, 5, 6, 7, 8\}$, or $h(v_i) = h(v_j) - 1$, if $h \in \{2, 4\}$. If the group of v_i is a central group, we increase by one unit the height of all the active vertices of the other central groups. Notice that the heights of the vertices of group h, for $h \leqslant 4$, are in the range $[(h-1) \cdot n + 1, h \cdot n]$, while the heights of the remaining vertices are greater than $4n$.

We now prove that Λ_i satisfies invariant I1 by showing that the addition of v_i does not destroy any existing visibility and that $\Lambda_i(v_i)$ sees all the other active vertices inside the appropriate regions. We have different cases depending on the group h of v_i.

– $h = 1$ **or** $h = 2$**.** The box $\Lambda_i(v_i)$ only intersects the regions $\alpha_{h',2}$, with $h' \neq 2$. Thus, the only visibilities that could be destroyed are those inside these regions. The visibilities in the regions $\alpha_{3,2}$, $\alpha_{4,2}$, $\alpha_{5,2}$, $\alpha_{6,2}$, $\alpha_{7,2}$, and $\alpha_{8,2}$ are not destroyed by the addition of v_i because the boxes representing the vertices of group 2 are taller than the box representing v_i and so are the boxes of any group h' with $h' > 2$. The existing visibilities in the region $\alpha_{1,2}$ are not destroyed because $r_{h,i}$ is short enough (in the x-direction) so that the existing boxes of groups 1 and 2 can still see each other in region $\alpha_{1,2}$. So, no visibility is destroyed for the vertices of group 2. The box $\Lambda_i(v_i)$ sees the box of the active vertex of group 1 or 2 via a ground visibility in region $\alpha_{1,2}$ and it sees the boxes of all the other active vertices inside the region $\alpha_{h',1}$, with $h' > 2$, above the boxes of group 1 (which are all shorter than it).

– $h = 3$ **or** $h = 4$**.** The proof of this case can be found in Appendix B.

– $h = 5$ **or** $h = 6$**.** The box $\Lambda_i(v_i)$ only intersect the regions $\alpha_{h,h'}$, with $h' \in \{5, 6, 7, 8\}$ and $h' \neq h$. However, it does not intersect any existing visibility inside these regions and therefore the addition of $\Lambda_i(v_i)$ does not destroy any existing visibility. The box $\Lambda_i(v_i)$ sees the active vertices of groups 1 and 2 inside $\alpha_{h,k}$ (with $h = 5$ or 6, and $k = 1, 2$) and above the boxes of group 1. The active vertices of groups 3 and 4 are seen inside $\alpha_{h,k}$ (with $h = 5$ or 6, and $k = 3, 4$) and above the boxes of group 3. Finally, the active vertices of the central groups are seen inside $\alpha_{h,k}$ (with $h = 5$ or 6, and $k > 4$) and above the boxes of group h. Recall that the active vertices of the central groups have been raised to have the same height as $\Lambda_i(v_i)$ (which is larger than the height of any other box in the central groups).

– $h = 7$ **or** $h = 8$**.** The proof of this case can be found in Appendix B.

The above construction can be done in $O(n)$ time. Since the normalized path decomposition can be computed in $O(n)$ time, the time complexity follows. □

5 2.5D Grid Box Representations

Next we give a tight bound on the edge density of graphs admitting a 2.5D-GBR. The proof, which appears in Appendix C, is based on the fact that a set of aligned (unit square) boxes induces an outerplanar graph. A square grid of boxes gives the bound.

Theorem 4. *Every n-vertex graph that admits a 2.5D-GBR has at most $4n - 6\sqrt{n}$ edges, and this bound is tight.*

In the next theorem we prove that deciding whether a given graph admits a 2.5D-GBR with a given footprint is NP-complete. We call this problem 2.5D-GBR-WITH-GIVEN-FOOTPRINT (2.5GBR-WGF). The reduction is from HAMILTONIAN-PATH-FOR-CUBIC-GRAPHS (HPCG), which is the problem of deciding whether a given cubic graph admits a Hamiltonian path [2].

Theorem 5. *Deciding whether a given graph G admits a 2.5D-GBR with a given footprint is NP-complete, even if G is a path.*

Proof sketch: We first prove that 2.5GBR-WGF is in NP. A candidate solution consists of a mapping of the vertices of G to the squares of the given footprint and a choice of the heights of the boxes. By Lemma 1 we can assign to each box an integer height in the set $\{1, 2, \ldots, n\}$. Thus the size of a candidate solution is polynomial in the size of the input graph. Given a candidate solution, we can test in polynomial time whether all edges of G are realized as visibilities. Thus, the problem is in NP.

We now describe a reduction from the HPCG problem. Let G_H be an instance of the HPCG problem, i.e. a cubic graph, with n_H vertices and m_H edges. We compute an orthogonal grid drawing Γ_H of G_H such that every edge has exactly one bend and no two vertices share the same x- or y-coordinate. Such a drawing always exists and can be computed in polynomial time with the algorithm by Bruckdorfer et al. [10]. We now use Γ_H as a trace to construct an instance $\langle G, F\rangle$ of the 2.5GBR-WGF problem, where G is a path and F a footprint, i.e, a set of squares. G is a path with $4n_H + m_H$ vertices and therefore F will contain $4n_H + m_H$ squares. The footprint F is constructed as follows. Γ_H is scaled up by a factor of four. In this way, every two vertices/bends are separated by at least four grid units. Each vertex v of Γ_H is replaced by a set $S(v)$ of four unit squares. In particular if vertex v has coordinates $(4x, 4y)$ in Γ_H, then it is replaced by the following four unit squares: $S_1(v)$ whose bottom-right corner has coordinates $(4x, 4y)$, $S_2(v)$ whose bottom-right corner has coordinates $(4x + 2, 4y)$, $S_3(v)$ whose bottom-right corner has coordinates $(4x, 4y-2)$, and $S_4(v)$ whose bottom-right corner has coordinates $(4x + 2, 4y - 2)$. We associate with each edge e

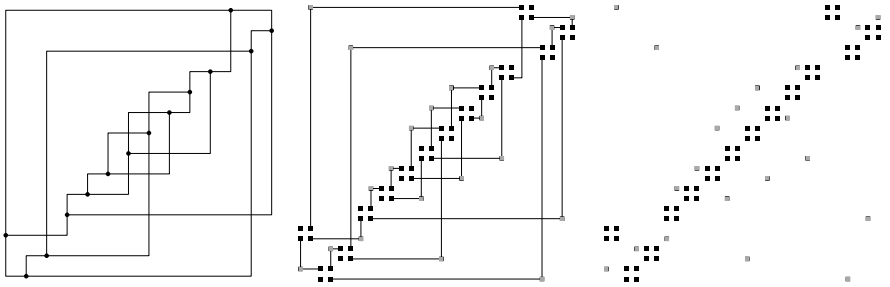

Fig. 4: (a) An orthogonal drawing of a cubic graph. (b) Construction of the footprint. Black (gray) squares are vertex (edge) squares. (c) The constructed footprint.

incident to a vertex v, one of the four squares in $S(v)$. If e enters v from West, North, South, or East, the square associated with e is $S_1(v)$, $S_2(v)$, $S_3(v)$, or $S_4(v)$, respectively. Let (u, v) be an edge of Γ_H and let $S_i(u)$ and $S_j(v)$ ($1 \leqslant i, j \leqslant 4$) be the squares associated with (u, v). The bend of $e = (u, v)$ is replaced by a unit square S_e horizontally/vertically aligned with $S_i(u)$ and $S_j(v)$. The set of squares replacing the vertices of Γ_H, which will be called *vertex squares* in the following, together with the set of squares replacing the bends, which will be called *edge squares* in the following, form the footprint F. Figure 4 shows an orthogonal drawing of a cubic graph and the corresponding footprint F. Observe that the footprint F is such that any two squares are separated by at least one unit and in each row/column there are at most three squares. Let F^* be a graph with a vertex for each square in F and an edge between two squares if and only if the two squares are horizontally or vertically aligned. It can be proved that G_H admits a Hamiltonian path if and only if F^* contains a Hamiltonian path, see Appendix C.

Consider the instance $\langle G, F\rangle$ of the 2.5GBR-WGF problem, where G is a path. We prove that G admits a 2.5D-GBR with footprint F if and only if F^* admits a Hamiltonian path. Every graph that can be represented by a 2.5D-GBR with footprint F is a spanning subgraph of F^* (because F^* has all possible edges that can be realized as visibilities in a 2.5D-GBR with footprint F). Thus, if G admits a 2.5D-GBR with footprint F, then G is a Hamiltonian path of F^* (recall that G is a path). Suppose now that F^* has a Hamiltonian path H^*. We show that we can choose the heights of the squares in F so that the resulting boxes form a 2.5D-GBR of G. Recall that in each row/column of F there are at most three squares. If an edge connects two squares that are consecutive along a row or column, then any choice of the heights is fine. If an edge connects the first and the last square of a row/column, then the heights of these two squares must be larger than the height of the square in the middle. We assign the heights to one square per step, in the order in which they appear along H^*. We assign to the first square a height equal to the number of squares (i.e., $4n_H + m_H$). Let h be

the height assigned to the current square S and let S' be the next square along H^*. If S and S' are consecutive along a row/column then the height assigned to S' is h. If S and S' are the first and the last square of a row/column then the height assigned to S' is h. If S is the first/last square of a row/column and S' is the middle square of the same row/column, then the height assigned to S' is $h-1$. If S is the middle square of a row/column and S' is the first/last square of the same row/column, then the height assigned to S' is $h+1$. It is easy to see that all heights are positive and that if an edge connects the first and the last square of a row/column, then the heights of these two squares are greater than the height of the square in the middle. This concludes the proof that G admits a 2.5D-GBR with footprint F if and only if F^* admits a Hamiltonian path. Since F^* has a Hamiltonian path if and only if G_H has a Hamiltonian path, G admits a 2.5D-GBR with footprint F if and only if G_H has a Hamiltonian path, which implies that the 2.5GBR-WGF problem is NP-hard. □

6 Open Problems

There are several possible directions for further study of 2.5D-BRs. Among them: (*i*) Study the complexity of deciding if a given graph admits a 2.5D-BR. We remark that deciding if a graph admits an RVR is NP-hard. (*ii*) Investigate other classes of graphs that admit a 2.5D-BR. For example, do 1-planar graphs or partial 5-trees always admit a 2.5D-BR? We remark that there are both 1-planar graphs and partial 5-trees not admitting an RVR. (*iii*) Study the 2.5D-BRs under the strong visibility model. For example, which bipartite graphs admit a strong 2.5D-BR?

References

1. M. E. Ahmed, A. B. Yusuf, and M. Z. H. Polin. Bar 1-visibility representation of optimal 1-planar graph. In *Electrical Inf. and Comm. Technology (EICT), 2013*, pages 1–5, 2014.
2. T. Akiyama, T. Nishizeki, and N. Saito. NP-completeness of the hamiltonian cycle problem for bipartite graphs. *J. of Inf. Processing*, 3(2):73–76, 1980.
3. F. Bernhart and P. C. Kainen. The book thickness of a graph. *J. of Combinatorial Theory, Series B*, 27:320–331, 1979.
4. Therese Biedl, Giuseppe Liotta, and Fabrizio Montecchiani. On visibility representations of non-planar graphs. In Sándor Fekete and Anna Lubiw, editors, *SoCG 2016*, volume 51 of *LIPIcs*, pages 19:1–19:16. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
5. H. L. Bodlaender. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Computing*, 25(6):1305–1317, 1996.
6. H. L. Bodlaender and T. Kloks. Efficient and constructive algorithms for the pathwidth and treewidth of graphs. *J. of Algorithms*, 21(2):358–402, 1996.
7. P. Bose, A. M. Dean, J. P. Hutchinson, and T. C. Shermer. On rectangle visibility graphs. In S. North, editor, *Proc. Symp. on Graph Drawing, GD '96*, volume 1190 of *LNCS*, pages 25–44, 1996.

8. P. Bose, H. Everett, S. P. Fekete, M. E. Houle, A. Lubiw, H. Meijer, K. Romanik, G. Rote, T. C. Shermer, S. Whitesides, and C. Zelle. A visibility representation for graphs in three dimensions. *J. Graph Algorithms and Applications*, 2(3):1–16, 1998.
9. F. Brandenburg. 1-visibility representations of 1-planar graphs. *J. of Graph Algorithms and Applications*, 18(3):421–438, 2014.
10. T. Bruckdorfer, M. Kaufmann, and F. Montecchiani. 1-bend orthogonal partial edge drawing. *J. of Graph Algorithms and Applications*, 18(1):111–131, 2014.
11. Paz Carmi, Eran Friedman, and Matthew J. Katz. Spiderman graph: Visibility in urban regions. *Computational Geometry*, 48(3):251 – 259, 2015.
12. A. M. Dean, J. A. Ellis-Monaghan, S. Hamilton, and G. Pangborn. Unit rectangle visibility graphs. *Electr. J. Comb.*, 15(1), 2008.
13. A. M. Dean, W. Evans, E. Gethner, J. D. Laison, Md. A. Safari, and W. T. Trotter. Bar k-visibility graphs. *J. of Graph Algorithms and Applications*, 11(1):45–59, 2007.
14. A. M. Dean and J. P. Hutchinson. Rectangle-visibility representations of bipartite graphs. *Disc. App. Math.*, 75(1):9–25, 1997.
15. A. M. Dean and J. P. Hutchinson. Rectangle-visibility layouts of unions and products of trees. *J. of Graph Algorithms and Applications*, 2(8):1–21, 1998.
16. A. M. Dean and N. Veytsel. Unit bar-visibility graphs. *Congressus Numerantium*, 160:161–175, 2003.
17. E. Di Giacomo, W. Didimo, W. S. Evans, G. Liotta, H. Meijer, F. Montecchiani, and S. K. Wismath. Ortho-polygon visibility representations of embedded graphs. *CoRR*, abs/1604.08797, 2016.
18. P. Duchet, Y. Hamidoune, M. Las Vergnas, and H. Meyniel. Representing a planar graph by vertical lines joining different levels. *Discrete Math.*, 46(3):319–321, 1983.
19. W. Evans, M. Kaufmann, W. Lenhart, T. Mchedlidze, and S. Wismath. Bar 1-visibility graphs and their relation to other nearly planar graphs. *J. of Graph Algorithms and Applications*, 18(5):721–739, 2014.
20. S. P. Fekete, M. E. Houle, and S. Whitesides. New results on a visibility representation of graphs in 3D. In F. Brandenburg, editor, *Graph Drawing*, volume 1027 of *LNCS*, pages 234–241. Springer Berlin Heidelberg, 1995.
21. S. P. Fekete and H. Meijer. Rectangle and box visibility graphs in 3D. *Int. J. Comput. Geometry Appl.*, 9(1):1–28, 1999.
22. S. Felsner and M. Massow. Parameters of bar k-visibility graphs. *J. of Graph Algorithms and Applications*, 12(1):5–27, 2008.
23. M. R. Garey, D. S. Johnson, and H. C. So. An application of graph coloring to printed circuit testing. *IEEE Trans. on Circuits and Systems*, CAS-23(10):591–599, 1976.
24. A. Gupta, N. Nishimura, A. Proskurowski, and P. Ragde. Embeddings of k-connected graphs of pathwidth k. *Disc. Applied Mathematics*, 145(2):242–265, 2005.
25. J. P. Hutchinson, T. Shermer, and A. Vince. On representations of some thickness-two graphs. *Comp. Geometry*, 13(3):161–171, 1999.
26. G. Kant, G. Liotta, R. Tamassia, and I. G. Tollis. Area requirement of visibility representations of trees. *Inf. Process. Lett.*, 62(2):81–88, 1997.
27. J. D. Kleitman, A. Gyárfás, and G. Tóth. Convex sets in the plane with three of every four meeting. *Combinatorica*, 21(2):221–232, 2001.
28. G. Liotta and F. Montecchiani. L-visibility drawings of IC-planar graphs. *Inf. Process. Lett.*, 116(3):217–222, 2016.

29. R. H. J. M. Otten and J. G. Van Wijk. Graph representations in interactive layout design. In *IEEE ISCSS*, pages 91–918. IEEE, 1978.
30. P. Rosenstiehl and R. E. Tarjan. Rectilinear planar layouts and bipolar orientations of planar graphs. *Discr. & Comput. Geom.*, 1:343–353, 1986.
31. T. C. Shermer. On rectangle visibility graphs III. external visibility and complexity. In *Cdn. Conf. on Comp. Geometry*, pages 234–239, 1996.
32. J. Štola. Unimaximal sequences of pairs in rectangle visibility drawing. In I. G. Tollis and M. Patrignani, editors, *Graph Drawing 2008*, volume 5417 of *LNCS*, pages 61–66. Springer Berlin Heidelberg, 2009.
33. I. Streinu and S. Whitesides. Rectangle visibility graphs: Characterization, construction, and compaction. In H. Alt and M. Habib, editors, *STACS 2003*, volume 2607 of *LNCS*, pages 26–37. Springer, 2003.
34. S. Sultana, Md. S. Rahman, A. Roy, and S. Tairin. Bar 1-visibility drawings of 1-planar graphs. In P. Gupta and C. Zaroliagis, editors, *Proc. Applied Algorithms: First Int. Conf. ICAA 2014*, pages 62–76. Springer Int. Publishing, 2014.
35. R. Tamassia and I. G. Tollis. A unified approach to visibility representations of planar graphs. *Discr. & Comput. Geom.*, 1(1):321–341, 1986.
36. R. Tamassia and I. G. Tollis. Representations of graphs on a cylinder. *SIAM J. Disc. Mathematics*, 4(1):139–149, 1991.
37. C. Thomassen. Plane representations of graphs. In *Progress in Graph Theory*, pages 43–69. AP, 1984.
38. S. K. Wismath. Characterizing bar line-of-sight graphs. In *Proc. 1st Symp. on Comp. Geometry*, pages 147–152, 1985.

Appendix

A Additional Material for Section 3

Proof of Lemma 1. By hypothesis G admits a 2.5D-BR Γ'. If every box of Γ' has a distinct integer height in the range $[1, n]$, the statement is true. If not, we can change the heights so to achieve this condition. Namely, denote by $b_1, b_2, \ldots, b_n$ the boxes of Γ' in non-decreasing order of height; we change the height of b_i to be i (for $i = 1, 2, \ldots, n$). Let Γ be the resulting representation and denote by $h'(b_i)$ the height of b_i in Γ' and by $h(b_i)$ the height of b_i in Γ. For any two boxes b_i and b_j, $h(b_i) < h(b_j)$ if and only if $h'(b_i) \leqslant h'(b_j)$, which means that no visibility has been destroyed by our change of the heights (while some new visibility may have been created). □

Proof of Lemma 2. For a given arrangement $\mathcal{R}$, choose ℓ_v and ℓ_h to be a vertical and horizontal line whose union intersects the maximum number of rectangles in $\mathcal{R}$. Suppose, for the sake of contradiction, that some rectangle $a \in \mathcal{R}$ is not intersected by $\ell_v \cup \ell_h$. Choose ℓ_v and ℓ_h so that they are closest to a without changing the set of rectangles intersected by their union. Assume w.l.o.g. that a lies in the positive quadrant of $\ell_v \cup \ell_h$. Let b be a rectangle that prevents ℓ_h from moving closer to a, that is, $b \cap \ell_h \neq \emptyset$ but $b \cap (\ell_v \cup \ell'_h) = \emptyset$, where ℓ'_h is ℓ_h translated in the $+y$ direction by any arbitrarily small positive amount. Let c be a rectangle that prevents ℓ_v from moving closer to A, that is, $c \cap \ell_v \neq \emptyset$ but $c \cap (\ell'_v \cup \ell_h) = \emptyset$, where ℓ'_v is ℓ_v translated in the $+x$ direction by any arbitrarily small positive amount. The line ℓ_h separates a and b, so $y(a) \cap y(b) = \emptyset$, which implies $x(a) \cap x(b) \neq \emptyset$. Similarly, using line ℓ_v, $y(a) \cap y(c) \neq \emptyset$.

By the conditions of the lemma, either $y(b) \cap y(c) \neq \emptyset$ or $x(b) \cap x(c) \neq \emptyset$. Suppose that $y(b) \cap y(c) \neq \emptyset$. Since $y(c)$ has non-empty intersection with both $y(a)$ and $y(b)$, any horizontal line that separates a and b must intersect c. Thus ℓ_h intersects c and $c \cap (\ell'_v \cup \ell_h) \neq \emptyset$ for all vertical lines ℓ'_v; a contradiction with the fact that c prevents ℓ_v from moving closer to a. We obtain a similar contradiction if $x(b) \cap x(c) \neq \emptyset$. □

Lemma 7. *K_{19} admits a 2.5D-BR.*

Proof. Let Λ be the box collection shown in Figure 2, where the footprint Λ_0 is depicted by a 2D drawing, while the heights of boxes are indicated by integer labels. We prove the statement by showing that Λ is a 2.5D-BR of K_{19}. We preliminarily observe that Λ_0 satisfies the necessary condition NC. The boxes of Λ are partitioned into five subsets T, R, L, B, C such that $|T| = |R| = 4$, $|L| = |B| = 3$ and $|C| = 5$; this partitioning is shown by the five rectangles with thick sides. For $S \in \{T, R, L, B\}$, it is easy to see that (for any assignment of heights to the boxes in Λ) $\Lambda(S)$ is a 2.5D-BR of $K_{|S|}$ in Λ, since any two boxes of $\Lambda(S)$ are ground visible in Λ_0. Hence, except for C, the intra-partition visibilities are ensured regardless of the heights of the boxes. We now show that even the inter-partition visibilities and the intra-partition visibilities of $\Lambda(C)$

exist if the heights of boxes are chosen as shown in Figure 2. We follow an incremental strategy in which, at each step, we add one of R, L, B, C, in order, to the current set of vertices, which is initialized to T.

Step 1: addition of R. Box $\Lambda(t_i)$ $(1 \leqslant i < |T|)$ may obstruct the visibility between a box in $\Lambda(R)$ and a box $\Lambda(t_{i'})$, with $i' > i$, in $\Lambda(T)$. This obstruction can be avoided, however, if (i) $\Lambda(T)$ has a staircase layout, i.e. the height of its boxes increases as $\Lambda(R)$ gets farther, and (ii) every box in $\Lambda(R)$ is not lower than any box in $\Lambda(T)$. Therefore, $\Lambda(T \cup R)$ is a 2.5D-BR of K_8 if $h(t_i) = i$ $(1 \leqslant i \leqslant |T|)$ and, for each $1 \leqslant j \leqslant |R|$, $h(r_j) \geqslant h_{max}(T)$, where $h_{max}(T)$ is the maximum box height in $\Lambda(T)$.

Step 2: addition of L. Consider the subset $S' = L \cup \{r_1\}$. As for the previous partite sets, $\Lambda(S')$ is a 2.5D-BR of $K_{|S'|}$ in Λ for any assignment of heights. However, box $\Lambda(r_i)$ $(1 \leqslant i < |R|)$ may prevent the (inter-partition) visibility between a box in $\Lambda(L)$ and a box $\Lambda(r_{i'})$, with $i' > i$, in $\Lambda(R)$. As before, these visibilities can be ensured if $\Lambda(R)$ has a staircase layout and every box in $\Lambda(L)$ is not lower than any box in $\Lambda(R)$. In particular, a possible assignment of heights is the following: $h(r_i) = h_{max}(T) + i - 1$ $(1 \leqslant i \leqslant |R|)$ and, for each $1 \leqslant j \leqslant |L|$, $h(l_j) \geqslant h_{max}(R)$. This assignment also implies the visibility between every box in $\Lambda(L)$ and every box in $\Lambda(T)$.

Step 3: addition of B. Consider now the subset $S'' = B \cup \{l_1\}$. As before, $\Lambda(S'')$ is a 2.5D-BR of $K_{|S''|}$ in Λ, independently from the choice of the heights of boxes. Furthermore, the inter-partition visibilities between $\Lambda(B)$ and $\Lambda(L)$ can be satisfied if $h(l_i) = h_{max}(R) + i - 1$ $(1 \leqslant i \leqslant |L|)$ and, for each $1 \leqslant j \leqslant |B|$, $h(b_j) \geqslant h_{max}(L)$. With this assignment of heights, the inter-partition visibilities between $\Lambda(B)$ and $\Lambda(T)$ and between $\Lambda(B)$ and $\Lambda(R)$ are ensured.

Step 4: addition of C. According to Fig. 2, every box $\Lambda(c_i)$ $(1 \leqslant i \leqslant 5)$ can see every other box $\Lambda(v)$ with $v \in T \cup R \cup L \cup B$, provided that $h(c_i) \geqslant h(v)$. Indeed, $\Lambda(c_i)$ is ground visible to any box in $\{\Lambda(t_1), \Lambda(r_1), \Lambda(b_1), \Lambda(l_1)\}$ and $\Lambda(T)$, $\Lambda(R)$, $\Lambda(L)$, $\Lambda(B)$ have a staircase layout with increasing box height as $\Lambda(C)$ gets farther. Therefore, if $h(c_i) \geqslant h_{max}(B)$ $(1 \leqslant i \leqslant 5)$, then all the inter-partition visibilities in Λ are satisfied. It remains to consider the intra-partition visibilities in $\Lambda(C)$. In this regard, the only visibility obstructions can be caused by $\Lambda(c_5)$, which may prevent the visibility between $\Lambda(c_1)$ and $\Lambda(c_3)$ ($\Lambda(c_2)$ and $\Lambda(c_3)$, respectively) if $h(c_1)$ and $h(c_3)$ ($h(c_2)$ and $h(c_4)$, respectively) are not both strictly greater than $h(c_5)$. Therefore, by choosing $h(c_5) = h_{max}(B)$ and $h(c_i) = h(c_5) + 1$ $(1 \leqslant i \leqslant 4)$, the intra-partition visibilities in $\Lambda(C)$ are satisfied, from which it follows that Λ is a 2.5D-BR of K_{19}. □

B Additional Material for Section 4

A *path decomposition* P of a graph $G = (V, E)$ is a sequence $P_1, \ldots, P_k$ of subsets of V, called *bags*, such that the following three properties hold:

- For every vertex u of G, there is a bag P_i (with $1 \leqslant i \leqslant k$) such that $u \in P_i$;
- For every edge (u, v) of G, there is a bag P_i (with $1 \leqslant i \leqslant k$) such that $u, v \in P_i$;

– For every vertex u, there exists two indices $1 \leqslant j \leqslant h \leqslant k$, such that u is contained in all bags P_i such that $j \leqslant i \leqslant h$ and in no other bag.

Let P_i be the bag of P with maximum size. The *width* of the path decomposition P is $|P_i| - 1$. The *pathwidth* of a graph G is the minimum width of any path decomposition of G.

A path decomposition $P = P_1, \ldots, P_k$ of a graph G of pathwidth p is normalized if $|P_i| = p + 1$ for i odd, $|P_i| = p$ for i even, and $P_{i-1} \cap P_{i+1} = P_i$ for i even.

For a fixed p, path decomposition of graphs with pathwidth p can be found in linear time [6, 5]. Given a path decomposition, a normalized path decomposition of the same width can be found in linear time [24].

Missing cases of the proof of Theorem 3. To complete the proof of Theorem 3, we need to prove the following cases.

$h = 3$ **or** $h = 4$**.** The box $\Lambda_i(v_i)$ only intersects the regions $\alpha_{4,k'}$, with $k' < 4$, and $\alpha_{k'',4}$, with $k'' > 4$. Thus, the only visibilities that could be destroyed are those inside these regions. The visibilities inside $\alpha_{4,1}$ and $\alpha_{4,2}$ are not destroyed because $r_{h,i}$ is placed so that the existing boxes of groups 4 can still see (to the left of $r_{h,i}$) the boxes of group 1 and 2 inside $\alpha_{4,1}$ and $\alpha_{4,2}$. The existing visibilities between boxes of group 4 and the boxes of group 3 are not destroyed because $r_{h,i}$ is short enough (in the x-direction) so that the existing boxes of groups 3 and 4 can still see each other (to the right of $r_{h,i}$) inside $\alpha_{4,3}$. The visibilities between the vertices of group 4 and vertices of group h', with $h' > 4$, are not destroyed because the boxes representing the vertices of group 4 are taller than $\Lambda_i(v_i)$ and so are the boxes of any group h' with $h' > 4$. So, no visibility is destroyed for the vertices of group 4. The box $\Lambda_i(v_i)$ sees the active vertex of group 1 with a visibility that is inside $\alpha_{3,1}$ or $\alpha_{4,1}$ and above the boxes of group 1 corresponding to non-active vertices (these boxes are shorter than the box of the active vertex of group one). Similarly, $\Lambda_i(v_i)$ sees the active vertex of group 2 with a visibility that is inside $\alpha_{3,2}$ or $\alpha_{4,2}$ and above the boxes of group 1 (including the active one). The box $\Lambda_i(v_i)$ sees the active vertex of group 3 or 4 via a ground visibility inside $\alpha_{4,3}$ and it sees the boxes of all the other active vertices inside $\alpha_{h,k}$ (with $h = 3$ or 4, and $k > 4$) above the boxes of group 3 (which are all shorter than it).

$h = 7$ **or** $h = 8$**.** The box $\Lambda_i(v_i)$ only intersects the regions $\alpha_{7,5}$, $\alpha_{8,5}$, and $\alpha_{8,7}$. Thus, the only visibilities that could be destroyed are those inside these regions. The visibilities between the vertices of groups 5 and 7 are not destroyed because $r_{h,i}$ is placed so that the existing boxes of groups 7 can still see the boxes of group 5 inside $\alpha_{7,5}$ above $r_{h,i}$ (in the y-direction). Similarly, the visibilities between the vertices of group 5 and the vertices of group 8 are not destroyed because the existing boxes of groups 8 can still see the boxes of group 5 inside $\alpha_{8,5}$ and below $r_{h,i}$ (in the y-direction). The visibilities between the vertices of group 7 and the vertices of group 8 are not destroyed

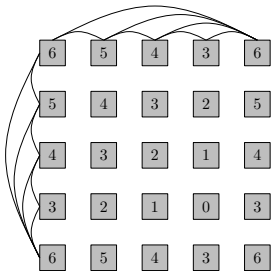

Fig. 5: A 2.5D-GBR with maximum edge density. The numbers in the square indicate the height of each box. The edges of the subgraph induced by the first row and the first column are shown.

because the existing boxes of groups 8 can still see the boxes of group 7 inside $\alpha_{8,7}$ and above $r_{h,i}$ (in the y-direction), if $h = 8$, or below $r_{h,i}$ (in the y-direction), if $h = 7$. The proof that v_i sees all the other active vertices is equal to the one in the case $h = 5$ or $h = 6$. □

C Additional Material for Section 5

Proof of Theorem 4. Consider a 2.5D-GBR Γ of an n-vertex graph G and let Γ_0 be the footprint of Γ. Γ_0 consists of a set of unit squares that are aligned along the columns and the rows of an integer grid. Let $B_1, B_2, \ldots, B_k$ be a set of boxes whose footprints share the same x-extent (i.e., are in the same column) or y-extent (i.e., are in the same row). Two boxes B_i and B_j, with $i + 1 < j$ can have a visibility only if they are both higher than any other box B_k with $i < k < j$. Thus, there cannot be four boxes B_i, B_j, B_k, B_h, with $i < j < k < h$ and such that there is a visibility between B_i and B_k and a visibility between B_j and B_h. If these two visibilities existed then B_j should be taller than B_k (in order to see B_h) and B_k should be taller than B_j (in order to see B_i). It follows that the subgraph G' of G that is represented by the boxes $B_1, B_2, \ldots, B_k$ has page number one and therefore is outerplanar[6]. This implies that the maximum number of edges of G' is $2k - 3$. Suppose that the unit squares in Γ_0 occupy R rows and C columns, and that row r_i (with $1 \leqslant i \leqslant R$) has n_{r_i} vertices while column c_i (with $1 \leqslant i \leqslant C$) has n_{c_i} vertices. The maximum number of edges in G is $\sum_{i=1}^{R}(2n_{r_i} - 3) + \sum_{i=1}^{C}(2n_{c_i} - 3) = 2n - 3R + 2n - 3C = 4n - 3(R + C)$. It is easy to see that this number is maximized when $R = C = \sqrt{n}$, i.e., when the squares of Γ_0 form a $\sqrt{n} \times \sqrt{n}$ grid. In this case the maximum number of edges is $4n - 6\sqrt{n}$.

For each $n = k^2$ a 2.5D-GBR that achieves the maximum edge density can be created by placing n boxes so that the footprint is a $k \times k$ grid and then choosing

[6] A graph $G = (V, E)$ has page number one if there exists a total order $\leqslant$ of V such that there are no two edges (u, v) and (w, z) with $u \leqslant w \leqslant v \leqslant z$. It is known that a graph has page number one if and only if it is outerplanar [3].

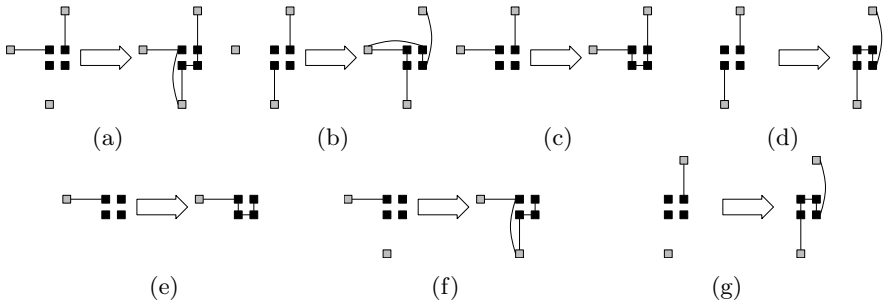

Fig. 6: Illustration for the proof of Theorem 5

.

the heights of the boxes along each row and column to form a decending sequence between two maxima. Figure 5 shows a 5×5 grid with appropriate heights indicated. The pattern is easily extended. □

Missing part of the proof of Theorem 5. To complete the proof of Theorem 5, we must prove that G_H admits a Hamiltonian path if and only if F^* contains a Hamiltonian path. Recall that F^* denotes a graph with a vertex for each square in F and an edge between two squares if and only if the two squares are horizontally or vertically aligned. Suppose first that G_H has a Hamiltonian path H. Each edge $e = (u, v)$ of H corresponds to two edges in F^*: one connecting a vertex square $S_i(u)$ $(1 \leqslant i \leqslant 4)$ to the edge square S_e, and one connecting S_e to a vertex square $S_j(v)$ $(1 \leqslant j \leqslant 4)$. Notice however that the set E_H of such edges does not form a Hamiltonian path of F^* because it does not contain all vertices of F^*. In particular, for each vertex v of G_H there are two vertex squares $S_i(v)$ and $S_j(v)$ $(1 \leqslant i, j \leqslant 4)$ that have no incident edge in E_H (for the end-vertices of H, the vertex squares without incident edges in E_H are three); also, the edge squares of all the edges that are not in H have no incident edges in E_H. We say that these edge squares are *orphans*. We now show that it is possible to select additional edges of F^* to create a Hamiltonian path. Each orphan edge square is assigned to one of the end-vertices of the edge corresponding to the square. The assignment is arbitrary, we only take care that at most one orphan edge square is assigned to each end-vertex of H. Let v be a vertex of G_H and suppose first that v is an internal vertex of H. We have two cases: an orphan edge square is associated with v or not. In both cases then we have two sub-cases: the vertex square of v with an incident edge of E_H are horizontally/vertically aligned or not. Figures 6(a)-6(d) shows for each sub-case how to select additional edges of F^* so that the four vertex squares of v and, possibly, the orphan edge square assigned to v are traversed by a simple path. The case when v is an end-vertex of H can be treated similarly, Figures 6(e)-6(g) shows the possible cases. It is easy to see that applying the transformations illustrated in Figure 6 to all the vertices v of G_H, we obtain a Hamiltonian path of F^*. Figure 7 shows a complete example.

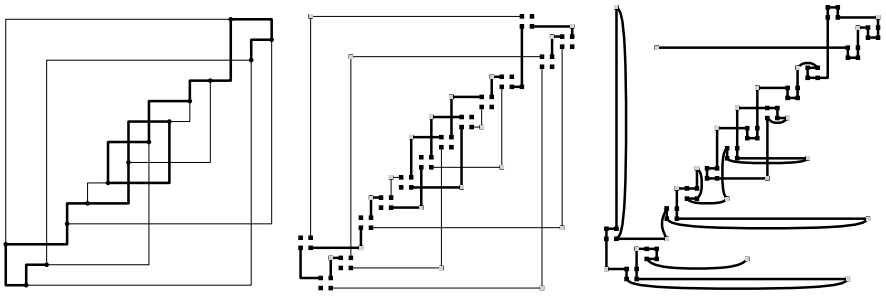

(a) (b) (c)

Fig. 7: (a) A Hamiltonian path H of the cubic graph G_H. (b) Constructing a Hamiltonian path of F^*. The bold edges are the edges of F^* that correspond to the edges of H. (c) A Hamiltonian path of F^*.

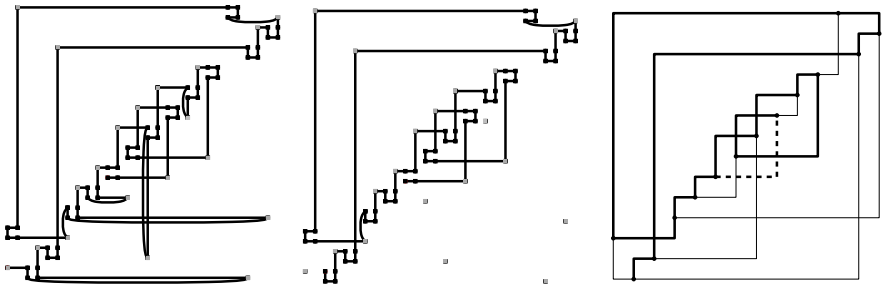

(a) (b) (c)

Fig. 8: (a) A Hamiltonian path H^* of F^*. (b) H^* simplified. (c) The candidate edges E_H of G_H. The dashed edge has to be removed to create a Hamiltonian path H of G_H.

Suppose now that F^* has a Hamiltonian path H^*. We show how to construct a Hamiltonian path H of G_H. We first make a simplification of H^*. Let S_e be an edge square of F^* and suppose that it is adjacent in H^* to two vertex squares $S_i(v)$ and $S_j(v)$ of the same vertex v. In this case S_e and its adjacent edges are removed from H^* and the edge connecting $S_i(v)$ and $S_j(v)$ is added to H^*. Analogously, if S_e is an end-vertex of H^*, we remove it (and its adjacent edge) from H^*. After this simplification, H^* is a simple path that traverses all vertex squares and a subset of the edge squares. Each edge square S_e that is still in H^* is adjacent to two vertex squares $S_i(u)$ and $S_j(v)$ of two different vertices u and v. Edge (u, v) of G_H is selected as a candidate edge for H. We now show that the set E_H of candidate edges forms a Hamiltonian path H of G_H, possibly after the removal of one or two edges. Let v be a vertex of G_H such that no end-vertex of H^* is a vertex square of v. We claim that the four vertex squares of v appear consecutively in H^*. Since G_H is a cubic graph, there can be at most

three edges of H^* with an end-vertex in $S(v)$ and the other end-vertex outside $S(v)$. However, since any other edge of H^* incident to a square of $S(v)$ must have the other end-vertex also in $S(v)$, if one or three such edges existed, then a square of $S(v)$ would be an end-vertex of H^*, but we are assuming that this is not the case. Thus, the four squares of $S(v)$ are connected to two squares not in $S(v)$ and therefore they must be consecutive in H^*. This implies that each vertex v of G_H, except at most two, has at most two incident candidate edges. The exceptions are the (at most) two vertices whose vertex squares include the end-vertices of H^*. This gives rise to two edges of E_H that can be removed to obtain a Hamiltonian path H. Figure 8 shows an example. □